



User Models, Intelligent Interface Agents, and Expert Systems,

UM 97 Workshop: Embedding User Models in Intelligent Applications
Sardinia , Italy, June 2-5 1997

Scott M. Brown, Robert A. Harrington, Eugene Santos Jr., and Sheila B. Banks

Department of Electrical & Computer Engineering
Air Force Institute of Technology
2950 P Street, Wright-Patterson AFB, OH 45433-7765

The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the US Government

Approved for public release; distribution unlimited.

User Models, Intelligent Interface Agents and Expert Systems

Scott M. Brown Robert A. Harrington Eugene Santos Jr. Sheila B. Banks

Department of Electrical and Computer Engineering, Air Force Institute of Technology
Wright-Patterson AFB, OH, USA

Voice: +1 937-255-3636 Fax: +1 937-656-4055
{sbrown, esantos, sbanks}@afit.af.mil

Abstract: The proliferation of user modeling as a means of accurately capturing the beliefs, desires, and intent of users is apparent. We present a dynamic, uncertainty-based user model knowledge representation for use in an intelligent interface agent called GESIA. GESIA's numerical uncertainty management representation, which has its roots in Bayesian networks, not only dynamically captures and models user behavior, but it also dynamically captures and models uncertainty in the agent's reasoning process. We critically discuss the construction of user models, relating their construction with user intent and the relevancy of observable properties in the user's "world." Our agent's usefulness is demonstrated in the domain of a generic expert system shell called PESKI.

1 Introduction

Nearly every application on the market today has some sort of user interface. As these applications become more complicated and complex, more is demanded of the user. The field of intelligent user interface research has risen to meet the challenge of making these user interfaces more usable. However, intelligent interface researchers have put little emphasis on improving the underlying structures that represent the intelligence of these interfaces.

It is widely agreed that basing decisions on an accurate cognitive model of the user is important for effective prediction of user intent (Opperman 1994, Thomas 1993). DeWitt (1993) notes that not all naturalistic (i.e., observable) properties play an interesting role in a user's causal model. That is, only certain observable actions and information in a user's "world" will have relevance to that user. Therefore, to effectively and efficiently capture user intent, a model should not attempt to model every possible action the user may exhibit, but only those that are relevant, i.e., most likely to be exhibited.

The underlying knowledge representation used in our agent is a Bayesian network¹ (Pearl 1988). Our architecture's knowledge representation is able to:

- provide a mathematically correct and semantically sound model for representing uncertainty and provide the means to show causal relationships between items;
- dynamically capture and model user behavior;
- dynamically capture and model uncertainty in the agent's reasoning process;
- provide unique user models for each user as well as user class models.

This paper is organized as follows. In Section 2, we present our application domain - a generic expert system shell. In Section 3, we present the intelligent interface agent GESIA, describing our knowledge representation. In Section 4, we critically discuss the construction

¹ See <http://www.afit.af.mil/Schools/EN/AI/> for an excellent resource on Bayesian networks.

of user models from the viewpoint of what is relevant to determine user intent. Finally, in Section 5 we present our results and conclusions.

2 Application Domain

Our research is part of a larger project exploring generic expert systems. PESKI (Probabilities, Expert Systems, Knowledge, and Inferencing)² is a generic expert system shell utilizing Bayesian knowledge bases (BKBs) as its knowledge representation (Santos & Banks 1995). PESKI is a well integrated tool, providing a user with knowledge acquisition, verification and validation, data mining, and inference engine tools. Furthermore, PESKI uses multiple communication modes, allowing a user to select the best possible way to view and input the information contained in the system. Currently, a structured text and a graphical communication mode exist. Future versions of PESKI will provide a natural language communication mode. Therefore, PESKI can be used by many types of users, from computer scientists and knowledge engineers with a lot of domain knowledge, to application users with (assumed) little domain knowledge. Due to the large amount of tools available to a user and the plethora of decisions that must be made, we desire to abstract out the difficulties and simplify the user's decisions. Our architecture accomplishes this by accurately modeling a user and assisting the user by suggesting useful actions the agent will take on the user's behalf. An example suggestion is shown in Figure 1. Figure 2 shows PESKI's verification and validation tool.



Figure 1. Tool and Communication Mode Suggestion

3 The GESIA Intelligent Interface Agent

GESIA (Generic Expert System Intelligent Assistant) is an intelligent interface agent conceived out of the development of a generic expert system shell (Harrington et al 1996a, 1996b). The goals of GESIA's development are twofold. The first goal is to provide useful and timely assistance to the user by accurately modeling the user. Secondly, we desire to

² See <http://www.afit.af.mil/Schools/EN/AI/PESKI> for a good description of PESKI.

maintain the domain independence of the agent. That is, we want to ensure the agent can easily be transplanted from one application domain to another.

The most important component of GESIA's architecture is the intelligent interface agent layer. This layer controls the communications and intelligence aspects of GESIA and is composed of three layers: the adaptation layer, the adaptive layer, and the communications layer. The adaptation layer manages and tracks all adaptations the user makes to the user interface. The adaptive layer communicates directly with the interface learning network gateway to perform interface initiated adaptations to the user interface based on perceived user behavior. Finally, the communications layer controls the various modes of communication available to the interface such as structured text, graphical manipulation, and natural language. The intelligent learning network is the heart of this layer. This knowledge representation captures, stores, and models user and interface behavior and has its roots in Bayesian networks.

4 Construction of User Models

Determining how to construct a network to accurately model a user is difficult at best. Furthermore, once we have built a user model, how do we modify it to more accurately model a user? Two main schools of thought exist on the construction of user models. The first uses "hand-coded" user models. That is, the system designer determines how best to model the users by constructing the user models a priori. See (Jameson 1995, Maes 1994) for examples. Hand-coded user models are typically static. Once they are designed, they will not change structure. The second method uses "machine-coded" user models. That is, a user model is constructed by the system as it "learns" more about the user. These models are dynamic (i.e., the structure changes over time). See (Jameson 1995, Harrington et al 1996a) for examples. Both methods have advantages and disadvantages and the particular domain will determine which method best meets the needs of the user model's construction. We, however, are concerned with other issues that impact both hand and machine-coded user models. These are discussed below.

4.1 Relevancy and User Intent

As mentioned previously, to accurately predict user intent, we must have an accurate cognitive model (i.e., user model). Modeling every possible naturalistic property in the user's world fortunately does not lead to the most accurate model. If this were not the case, we would have little hope in using numerical uncertainty management techniques such as Bayesian networks due to the computational inefficiency of large networks (Cooper 1990).

DeWitt (1993) uses the term "causally efficacious" to describe naturalistic properties that "play any interesting causal role in cognitive functions." He argues not everything observable is of interest when we make decisions. Dewitt's philosophical argument has direct analogy in user models. A user model must not include every possible piece of information about "the world." To use Dewitt's example, while although the manufacturer of my stereo, the amount of dust on top of it, and the weight of the unit are all observable properties, none of them would likely have any bearing on my reasoning as to why the stereo does not work. To include such information in a user model needlessly complicates the model, not only semantically, but

computationally. To take this analogy one step further, the less causally efficacious a property is to another, the more likely it is we can ignore it.



Figure 2. PESKI Verification and Validation Tool

We use the term relevancy set to describe those properties (i.e., nodes) included in a user’s interface learning network. The two terms are used interchangeably; however, the network fully captures the causality of the model (i.e., nodes, arcs, and probabilities), while the relevancy set only captures the nodes in the network. For hand-coded models, the relevancy set will not change. For machine-coded models, the relevancy set may change with use. A relevancy set are those properties (i.e., nodes) that are immediately causally efficacious to the decision under consideration. As an example, a user’s interface learning network may contain the possible communication modes and tools a user may use, given the user’s class and individual preferences learned by the user’s interface learning network, as well as several knowledge bases used previously. Therefore, the interface learning network presents a computationally efficient and semantically meaningful view of the user’s relevancy set at any given time. The concept of a relevancy set is a good tradeoff between computational complexity and representational exactness, while maintaining full semantics.

GESIA currently uses a dynamic “hand-coded” interface learning network (Brown 1997). We determine *a priori* the actions we will monitor. This is not unlike Maes’ hand-coded situations (Maes 1994). This a priori determination limits the number of user actions we must monitor in our system. While although most “hand-coded” user models are static, ours allows the dynamic addition and deletion of nodes as the intelligent interface agent learns more about the user. We limit the number of nodes allowed in the user’s interface learning network at any one time. If the user performs an action that is not represented in the current relevancy set

(i.e., the node is not in the network), we add it to the network. If we have reached our network size limitation, we delete the lowest probability node from the network. In this way, the most relevant (i.e., highest probability) actions (nodes) are in our relevancy set at any point in time.

4.2 User Classes

A user class is a generalization of a number of users sharing common characteristics. User classes have been used in a number of systems for various reasons (Jameson 1995, Maes 1994). PESKI segregates users into the following user classes: application users, application experts, knowledge engineers, and computer scientists (Harrington et al 1996a). User classes in PESKI serve two main purposes. New users to the system will have their individual interface learning nodes (i.e., Bayesian network random variables) set to the most uncertain probability (i.e., 0.50). However, a user's user class information learning nodes will contain probabilities of his class. Therefore, the user class helps bias the agent initially towards this class. As the user uses PESKI, his preferences are captured appropriately in the individual user interface learning nodes, and therefore, GESIA makes better predictions concerning the user's intent. Analogously, a real-world personal assistant can not be expected to accurately predict an employer's actions the first day. However, based on previous employers with similar backgrounds, the personal assistant attempts to determine what information her employer will need and when. As she learns more about her employer's behaviors, beliefs, and intentions, prediction becomes more accurate. Secondly, by collecting user class data, system designers can determine what communication modes, tools, etc. are being used by a particular class. This may allow for future improvements to the system based on usage.

Several problems can arise from user classification. First, a user may belong to more than one user class. For example, consider the computer scientist concerned with the development of the expert system (i.e., requirements, design, implementation, maintenance) who is also the knowledge engineer. As she uses the system, some of her actions are best categorized by the computer scientist class, while others are best characterized by the knowledge engineer user class. Both the user and the user class may suffer as a result. Another major problem is the misclassification of users. PESKI allows the user to select his user class upon the first login. If this user misclassifies himself, it is possible that his actions may be diametrically opposed to those of the class in which he has placed himself. Once again, both the user and the user class suffer. We desire to minimize these problems by identifying when they occur. Proposed solutions to these problems are discussed elsewhere (Brown 1997).

5 Results and Conclusions

We performed four usability studies to determine the usefulness of GESIA's user model in PESKI (Banks 1997). Our studies showed that GESIA reduced a user's workload substantially by correctly predicting the communication modes, tools, and knowledge bases to be used by the users. Users expressed that GESIA did slightly impact the responsiveness of the overall system. This was due to the inferencing scheme used in the intelligent learning network. A vastly improved inferencing scheme has been implemented and responsiveness has increased dramatically. Overall, users welcomed the addition of GESIA into the system and agreed that it accurately predicted their actions.

In this paper, we have presented a new domain for the use of Bayesian-based user models. Our architecture is not only capable of dynamically modeling user behavior, but also in learning when it is wrong. When designing user models, whether hand or machine-coded, we must realize that not everything that is observable is needed to make our predictions about user intent. The concept of a relevancy set for use in a user model is a good tradeoff between computational complexity and representational exactness. Furthermore, the semantics of the user model are maintained. The classification of users into meaningful groups can be a benefit to systems employing user models. User classes can help model a new user, before the system has collected meaningful information on the user. We feel that this new domain, integrating the fields of artificial intelligence, user modeling, human-computer interfaces, and philosophy will reap great benefits for users of the systems employing these sorts of techniques.

References

- Banks, S. B., Harrington R. A., Santos Jr., E., and Brown, S. M. (1997) Usability testing of an intelligent interface agent. To appear in *Proceedings of the Sixth International "Interfaces '97" Conference*.
- Brown, S. M., Harrington, R. A., Santos Jr., E., and Banks, S. B. (1997). A dynamic bayesian intelligent agent. To appear in *Proceedings of the Sixth International "Interfaces '97" Conference*.
- Cooper, G. F. (1990). The computational complexity of probabilistic inference using bayesian belief networks. *Artificial Intelligence* 42:393-405.
- DeWitt, R. (1993). Vagueness, semantics, and the language of thought. In *Psyche*, volume 1. December.
- Harrington, R. A., Banks, S., and Santos Jr., E. (1996). Development of an intelligent user interface for a generic expert system. In Michael Gasser, ed., *Online Proceedings of the Midwest Artificial Intelligence and Cognitive Science Conference*. URL: <http://www.cs.indiana.edu/event/maics96/Proceedings/harrington.html>.
- Harrington, R. A., Santos Jr., E., and Banks, S. (1996). GESIA: uncertainty-based reasoning for a generic expert system intelligent user interface, *International Conference on Tools in Artificial Intelligence 1996*.
- Jameson, A. (1995). Numeric uncertainty management in user and student modeling: An overview of systems and issues. *User Modeling and User-Adapted Interactions* 5.
- Santos Jr., E. and Banks, D. O. (1995). A probabilistic framework for representing uncertainty. Technical report, Air Force Institute of Technology, Wright-Patterson AFB OH USA.
- Maes, P. (1994). Agents that reduce work and information overload. *Communications of the ACM* 37(7):811-821, July.
- Opperman, R. (1994). Adaptively supported adaptivity. *International Journal of Human-Computer Studies* 40:455-472.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman.
- Thomas, C. G. (1993). Design, implementation, and evaluation of an adaptive user interface. *Knowledge-Based Systems* 6(4):230-238, December.